

MEANDERS IN A CAYLEY GRAPH

H. TRACY HALL

ABSTRACT. A meander of order n is a simple closed curve in the plane which intersects a horizontal line transversely at $2n$ points. (Meanders which differ by an isotopy of the line and plane are considered equivalent.)

Let Γ_n be the Cayley graph of the symmetric group S_n as generated by all $\binom{n}{2}$ transpositions. Let Λ_n be any interval of maximal length in Γ_n ; this graph is the Hasse diagram of the lattice of noncrossing partitions. The meanders of order n are in one-to-one correspondence with ordered pairs of maximally separated vertices of Λ_n .

1. INTRODUCTION

The lattice of noncrossing partitions and its variants have received a fair amount of attention in recent years as a crossing grounds of several seemingly unrelated disciplines of mathematics, including Artin groups, cluster algebras, and free probability. We here record a connection between noncrossing partitions and the set of meanders, or isotopy classes of simple closed curves transversely intersecting an infinite line in the plane. It is well recognized that meanders correspond to certain pairs of noncrossing partitions, and indeed attempts have been made to make use of this description to aid in the notoriously difficult meander enumeration problem. To our knowledge, however, it is a new observation that the pairs of noncrossing partitions which give rise to meanders are exactly those whose path distance is maximal in the Hasse diagram of the lattice of noncrossing partitions—which graph can itself be regarded as a geometric interval in the Cayley graph of the symmetric group generated by all two-element cycles. It remains to be seen whether this is any aid to the enumeration problem; the difficulty seems to stem from the fact that the Hasse diagram is properly a directed graph, and minimal paths connecting elements can “change direction” an arbitrary number of times.

In Section 2 we define meanders, noncrossing matchings and partitions, and the Cayley graph Γ_n and lattice Λ_n . In Section 3 we show that Λ_n is a geometric interval in Γ_n and that meanders correspond to pairs of vertices in Λ_n whose path-distance is maximal, and finish with some unanswered questions arising from this bijection.

The author wishes to thank Reinhard Franz, who first introduced me to the problem of meander enumeration, as well as organizers of combinatorics seminars at Brigham Young University and the University of California, Berkeley, where these results were first presented several years ago. The 2005 Workshop “Braid Groups, Clusters and Free Probability” at the American Institute of Mathematics was valuable not only as a source of new information and new questions but also as an encouragement to write up older answers and make them known to the broader mathematical community. Further encouragement came from David Savitt, who recently observed and demonstrated independently [Sav06] the main result: that




FIGURE 1. The eight meanders of order 3.

the number of components of a system of meanders is measured exactly by path distances in the lattice of noncrossing partitions.

2. PRELIMINARIES

2.1. Meanders. A meander of order n is a simple closed curve in the plane which intersects an infinite horizontal line transversely in $2n$ points, which points we label $\bullet 1, 1\bullet, \bullet 2, 2\bullet, \dots, \bullet n, n\bullet$. (Two meanders which differ only by isotopies of the line and the upper and lower half-planes are considered equivalent, making meanders combinatorial objects.) For example, Figure 1 depicts the eight possible meanders of order 3. We will refer to the collection of points $\{\bullet 1, \bullet 2, \dots, \bullet n\}$ as *odd* points, and the points $\{1\bullet, 2\bullet, \dots, n\bullet\}$ we will call *even*. Meanders have a long history and still await a satisfactory enumeration; see for example [LZ93].

A meander is characterized by the labels of the endpoints of arcs both above and below the dividing line, and since every arc encloses either both endpoints or neither endpoint of every other arc, each arc encloses an even total number of points other than its own endpoints, and hence one of those endpoints is even and one is odd. There is thus a unique permutation σ of the integers $1, \dots, n$ such that $\bullet\sigma(i)$ is an endpoint of an arc in the upper half-plane if and only if the other endpoint of the arc is $i\bullet$, and similarly a unique permutation τ arising from the endpoints of arcs in the lower half-plane. The pair of permutations (σ, τ) is a complete invariant of meanders of a given order, but not every pair of permutations arises in this way.

Our principal concern is to characterize which pairs of permutations (σ, τ) define a meander, starting from topological considerations in the plane (namely: no pair of arcs may intersect, and the union of arcs must have a single component) to arrive at a characterization in terms of path lengths in a Cayley graph of the symmetric group.

2.2. Non-crossing partitions. We first consider the collection of arcs on a single side of the dividing line. Any permutation σ gives a matching of even points $i\bullet$ to odd points $\bullet\sigma(i)$, but for σ to represent half of a meander, there must exist disjoint arcs in a single half-plane which connect the pairs of matched points. Such a system of n arcs we call a *noncrossing matching* of order n . Every such system is canonically equivalent to a noncrossing partition, which we define as follows:

Definition 2.1. Let n points on the boundary of a circular disc be labeled 1 through n in counter-clockwise order, and let \sim be an equivalence relation on the set of marked points. The partition arising from the relation \sim is said to be *noncrossing*




FIGURE 2. The noncrossing partition of order 8 coming from the permutation $(1347)(56)$, with convex hulls in black. The 16 points $\bullet 1, 1\bullet$, etc. are paired by the boundary arcs of the shaded region.

if no two blocks of the partition “cross”—that is, if whenever two points $i \sim k$ and two other points $j \sim \ell$ have labels satisfying $i < j < k < \ell$, then $j \sim k$ as well and all four points are part of a single block. Equivalently, form the convex hull in the disc of each block of points, and the partition is noncrossing if and only if the convex hulls are disjoint.

Proposition 2.2. *The noncrossing matchings of order n in a single half-plane are in canonical bijection with noncrossing partitions of order n .*

Proof. To begin with, we take the one-point compactification of the half-plane (together with a reversal of orientation in the case of the lower half-plane) so that the arcs join points (marked $\bullet 1, 1\bullet, \dots, \bullet n, n\bullet$ in counter-clockwise order) on the boundary of a disc. These n arcs divide the disc into $n + 1$ connected components. We define a relation $i \sim j$ if the section of disc boundary which joins $\bullet i$ to $i\bullet$ is in the same connected component as the section joining $\bullet j$ to $j\bullet$; the relation \sim defines a noncrossing partition of order n . Conversely, if we start with a noncrossing partition of n points marked $1, \dots, n$ in counter-clockwise order around a disc, form convex hulls of the blocks of the partition, and take a regular neighborhood of the union of convex hulls, the boundary of the regular neighborhood will consist of n properly embedded disjoint arcs in the disc, whose endpoints surround the points

$\{1, 2, \dots, n\}$ in pairs $(\bullet 1, 1 \bullet), \dots, (\bullet n, n \bullet)$. (Figure 2 illustrates this correspondence for the example of the noncrossing partitions whose blocks are $\{1, 3, 4, 7\}$, $\{2\}$, $\{5, 6\}$, and $\{8\}$.) \square

2.3. The lattice Λ_n of noncrossing partitions. The set of all noncrossing partitions of a given order n we call Λ_n . This set has a great deal of structure and symmetry; in particular there is an operation which reverses the parity of the points $\bullet 1, 1 \bullet$, etc., and which induces a symmetry of order $2n$ on Λ_n . To every noncrossing partition we assign a dual as follows:


Definition 2.3. Let p be a noncrossing partition of order n corresponding to a permutation σ which pairs points $i \bullet$ with $\bullet \sigma(i)$. Then \vec{p} , which we call the *dual* of p , is the noncrossing partition arising from the noncrossing matching which pairs $\bullet i$ with $(\sigma(i) - 1) \bullet$ (where the subtraction is performed modulo n).

In other words, \vec{p} is the partition that arises if $1 \bullet$ is relabeled $\bullet 1$, $\bullet 2$ is relabeled $1 \bullet$, and so on cyclically (pulling the first point $\bullet 1$ around to become the last point $n \bullet$), so that the equivalence relation is given by the complementary connected components to those that defined p . Referring to Figure 2, we see for example that the partition dual to $\{\{1, 3, 4, 7\}, \{2\}, \{5, 6\}, \{8\}\}$ is $\{\{1, 2\}, \{3\}, \{4, 6\}, \{5\}, \{7, 8\}\}$. Note that the dual of the dual does not yield the same noncrossing partition back again, but rather a relabeling of the original by $1/n$ of a full rotation.

Non-crossing partitions have a natural partial order \preceq by refinement, where the “least” partition has n blocks and the “greatest” has only a single block containing all elements $1, \dots, n$. Duality reverses this partial order: $p \preceq q$ if and only if $\vec{q} \preceq \vec{p}$. Refinement \preceq gives Λ_n the structure of a lattice, where the join of two partitions comes from taking the union of the two equivalence relations and then extending by both transitivity and the noncrossing constraint—that is, if $i < j < k < \ell$ and $i \sim k, j \sim \ell$, we require $j \sim k$ as well. The meet operation is similarly well-defined (for example, as the anti-dual of the join of the duals). The lattice Λ_n is graded by n minus the number of blocks in the partition, and every maximal chain has the same length: q covers p in the partial order if and only if p is obtained by splitting a block of q into two blocks (in a noncrossing way), or in other words if $p \preceq q$ and they differ by 1 in the grading. Figure 3 illustrates the lattice of noncrossing partitions of order 4.

2.4. The Cayley graph Γ_n of transpositions in S_n . One place where noncrossing partitions have turned up prominently is in a particularly elegant solution to the word problem for the braid group of order n . Given any finite-type Artin group, there are two canonical Garside structures, one arising from the positive monoid on the standard generating set (the “standard” Garside structure), and one which is in many senses dual to it. In the case of the braid group B_n , the dual Garside monoid is generated by a set of $\binom{n}{2}$ conjugates of the standard generators, whose image in the corresponding Coxeter group S_n is the set of transpositions (those permutations that exchange two elements, leaving the others fixed).

We will consider permutations as vertices of the Cayley graph Γ_n of S_n generated by all transpositions. (Since these generators are involutions, we may take Γ_n to be a simple, undirected graph.) We write permutations in disjoint cycle notation, optionally omitting fixed points. We adopt the convention that permutations act on the left, so that for example $(12)(23) = (123)$. We give names to two distinguished vertices of Γ_n : the identity permutation e , and the successor function

FIGURE 3. The Hasse diagram of the lattice Λ_4 .

$s = (123 \dots n)$. Given any two permutations σ and τ , we define the distance $d(\sigma, \tau)$ in Γ_n as the minimal length of a path from σ to τ . By considering each edge of Γ_n to be a unit interval, we may extend d to the entire graph.

The properties of Γ_n as a Cayley graph ensure that d is symmetric and equivariant with respect to left and right multiplication, so that in particular $d(\sigma, \tau) = d(e, \sigma^{-1}\tau)$. We also have the following

Proposition 2.4. *If $\sigma \in S_n$ has r orbits (including fixed points), then $d(e, \sigma) = n - r$.*

Proof. Let $\sigma \in S_n$ be written as a product of disjoint cycles

$$(i_{11}i_{12} \dots i_{1\ell_1})(i_{21} \dots i_{2\ell_2}) \dots (i_{r1}i_{r2} \dots i_{r\ell_r}),$$

including fixed points, so that ℓ_j may equal 1 sometimes and so that $\sum_{j=1}^r \ell_j = n$. On the one hand we have

$$(i_{j1} i_{j2})(i_{j2} i_{j3}) \dots (i_{j\ell_j-1} i_{j\ell_j}) = (i_{j1} i_{j2} \dots i_{j\ell_j}),$$

so each cycle of length ℓ_j is a product of $\ell_j - 1$ transpositions, and

$$d(e, \sigma) \leq \sum_{j=1}^r (\ell_j - 1) = n - r.$$

But on the other hand multiplication of a permutation by a transposition cannot reduce the number of cycles by more than 1, so $d(e, \sigma) \geq n - r$ and we have equality. \square

The definition of duality for noncrossing partitions has a convenient translation in terms of permutations: If σ is the permutation such that $i \bullet$ is paired with $\bullet \sigma(i)$, then the permutation such that $\bullet i$ is paired with $(\sigma(i) - 1) \bullet$ is $\sigma^{-1}s$. To revisit Figure 2 for example, we have $(8)(65)(2)(7431)(12345678) = (12)(3)(46)(5)(78)$. Applying the dual twice gives us $(\sigma^{-1}s)^{-1}s = s^{-1}\sigma s$, or in other words conjugation by $1/n$ of a full rotation.

3. RESULTS

3.1. The lattice Λ_n of noncrossing partitions within Γ_n . Given a metric space (X, d) and a pair of points $x, z \in X$, the *interval* $\mathcal{I}_{x,z}$ from x to z is the union of all global geodesics between x and z , that is

$$\mathcal{I}_{x,z} = \{y \in X : d(x, y) + d(y, z) = d(x, z)\}.$$

The following theorem is implicit in the standard developments of the dual Gar- side structure on the braid group, although it is not usually stated in terms of an interval in the Cayley graph, nor is it usually proved in terms of the genus of a constructed surface.

Theorem 3.1. *Let each vertex in the Hasse diagram Λ_n of the lattice of noncrossing partitions be labeled with the permutation whose cycles are the blocks of the partition, each cycle written in numerically increasing order. Then this labeled Hasse diagram is the interval $\mathcal{I}_{e,s}$ in the metric space (Γ_n, d) .*

Proof. The conventions we have chosen—that the noncrossing partition corresponding to a permutation σ connects $i\bullet$ to $\bullet\sigma(i)$, and that the boundary points of a noncrossing partition are labeled in counter-clockwise order around the disc—ensure that the permutation σ associated to a noncrossing partition p consists of a cycle for each block of p , such that each cycle can be written in numerically increasing order.

If a permutation σ does come from a non-crossing partition, then the $n - 1$ connected components of the disc divided by the arcs of the noncrossing matching correspond to the cycles of σ and of the dual $\sigma^{-1}s$. Each connected component is a topological disc, and they are glued together along the arcs. We imitate this construction in the case of an arbitrary permutation, obtaining a disc exactly when the permutation lies in Λ_n .

Let σ be an arbitrary permutation, possibly crossing; we extend the definition of the dual $\vec{\sigma}$ as $\sigma^{-1}s$. We construct an oriented surface with boundary as follows: for each cycle $c = (c_1 c_2 \dots c_k)$ of σ , create a topological disc with $2k$ points marked $\bullet c_1, c_1\bullet, \bullet c_2, \dots, c_k\bullet$ in counterclockwise order, where of course $c_{i+1} = \sigma(c_i)$. The k edges from $c_i\bullet$ to $\bullet\sigma(c_i)$ we label c_i and orient counter-clockwise (that is, in the same order as previously listed) so that the union of discs for all cycles of σ has, for each label $1 \leq i \leq n$, exactly one edge with the label i going from the point marked $i\bullet$ to the point marked $\bullet\sigma(i)$. Similarly, for each cycle $c = (c_1 c_2 \dots c_k)$ of $\vec{\sigma}$, we mark $2k$ points on the boundary of a disc as $c_1\bullet, \bullet(c_1 + 1), c_2\bullet, \dots, \bullet(c_k + 1)$, where $c_i = \sigma^{-1}s(c_{i-1})$ and thus $c_{i-1} + 1 = \sigma(c_i)$. The k edges from $\bullet\sigma(c_i)$ to $c_i\bullet$ we label c_i and orient clockwise (that is, in the reverse order of that previously listed) so that the union of discs for all cycles of $\vec{\sigma}$ also has a full complement of edges labeled i and going from $i\bullet$ to $\bullet\sigma(i)$.

If σ represented a noncrossing partition to begin with, then what we have just described in rather intricate detail is the collection of connected components obtained by cutting a disc along the arcs of the corresponding noncrossing matching—refer for example to Figure 2. For an example of what happens in the more general case, refer to Figure 4 which illustrates the collection of marked discs obtained from the permutation $(15)(2436)$.

Identifying pairs of oriented edges labeled i between the two collections of discs, we obtain a connected oriented surface Σ_σ whose single boundary component has




FIGURE 4. The components of Σ_σ for $\sigma = (15)(2436)$ (top row from σ ; bottom row from $\vec{\sigma}$).

$2n$ vertices marked $\bullet 1, 1\bullet, \dots, \bullet n, n\bullet$ in order. The permutation σ represents a non-crossing partition if and only if Σ_σ is a disc, which is true (for a connected compact surface with one boundary component) if and only if the Euler characteristic of Σ_σ is 1.

We have $2n$ vertices for Σ_σ and $3n$ edges ($2n$ on the boundary and n in the interior). Each cycle of σ contributes one face, and since $d(e, \sigma) = n - k$ where k is the number of cycles, we have $n - d(e, \sigma)$ faces coming from σ . Similarly, we have $n - d(e, \sigma^{-1}s) = n - d(\sigma, s)$ faces coming from $\vec{\sigma}$. It follows that the Euler characteristic of Σ_σ is 1 exactly when

$$d(e, \sigma) + d(\sigma, s) = n - 1 = d(e, s)$$

and thus that the vertices of Λ_n are exactly those permutations which lie in $\mathcal{I}_{e,s}$. Indeed, the genus of Σ_σ is

$$\frac{1}{2} [d(e, \sigma) + d(\sigma, s) - d(e, s)]$$

and measures exactly the degree to which σ fails to lie in the interval.

It remains only to show that the edges of Λ_n are the same as those of $\mathcal{I}_{e,s}$. Every edge in $\mathcal{I}_{e,s}$ is part of a path of length $n - 1$ from e to s , and each step of such a path unites two orbits of a permutation; this is a covering relation in the refinement order \preceq of noncrossing partitions. Conversely, every covering relation in the order \preceq joins two partition blocks $(i_1 i_2 \dots i_j)$ and $(i_{j+1} \dots i_\ell)$ to form $(i_1 i_2 \dots i_\ell)$; this can be accomplished by left multiplication by the transposition $(i_1 i_{j+1})$, which decreases the distance to s in Γ_n and is thus an edge of $\mathcal{I}_{e,s}$. \square

Remark 3.2. Given any pair of permutations $\sigma, \tau \in \Gamma_n$ for which $d(\sigma, \tau) = n - 1$, $\sigma^{-1}\tau$ consists of a single cycle (conjugate to s) and $\mathcal{I}_{\sigma, \tau}$ is isomorphic to $\mathcal{I}_{e,s}$ (since the generating set of Γ_n is invariant with respect to conjugation). There are thus $n!(n - 1)!$ isometrically embedded copies of Λ_n in Γ_n .

Having defined Λ_n in terms of the metric d on Γ_n , it is natural to define a metric $d_\Lambda(\sigma, \tau)$ as the length of the shortest path from σ to τ which is contained entirely within Λ_n . (In fact, as we shall shortly see, $d_\Lambda = d$.)

3.2. Meanders in Λ_n . We now turn our attention to pairs of permutations (σ, τ) where each of the permutations σ, τ belongs to $\Lambda_n = \mathcal{I}_{e,s}$. When we construct the

noncrossing matching corresponding to σ in the upper half-plane and the (reflection of the) noncrossing matching corresponding to τ in the lower half-plane, we obtain a *system of meanders*, a collection of simple closed curves with some number of components $\Pi_0(\sigma, \tau)$. The meanders are exactly those pairs (σ, τ) for which $\Pi_0(\sigma, \tau) = 1$.

Theorem 3.3. *For any pair of permutations $\sigma, \tau \in \Lambda_n$, the following are equivalent:*

- (1) $\Pi_0(\sigma, \tau) = k$
- (2) $d(\sigma, \tau) = n - k$
- (3) $d_\Lambda(\sigma, \tau) = n - k$

Proof. Let M be the system of meanders corresponding to the pair (σ, τ) , where σ and τ are both permutations belonging to the lattice Λ_n . As usual, we label the transverse crossings of M as $\bullet 1, 1\bullet, \dots, \bullet n, n\bullet$ in order.

As each odd point $\bullet i$ is connected to an even point by an arc in the upper half-plane, in order to count the number of components $\Pi_0(\sigma, \tau)$ of the system of meanders M it suffices to find which components contain which of the even points $i\bullet$. We define a permutation of the integers $1, \dots, n$ as follows: starting at an even point $i\bullet$, first follow the arc connecting $i\bullet$ in the lower half-plane to $\bullet \tau(i)$, then follow the arc in the upper half-plane which connects $\bullet \tau(i)$ to $\sigma^{-1}\tau(i)\bullet$. Repeating this permutation we will eventually traverse the entire connected component of M containing $i\bullet$; it follows that the components of M are in one-to-one correspondence with the cycles of the permutation $\sigma^{-1}\tau$. We thus have

$$d(\sigma, \tau) = d(e, \sigma^{-1}\tau) = n - k,$$

where $k = \Pi_0(\sigma, \tau)$ is the number of connected components of M .

Now to show the equivalence of d and d_Λ . Every path in Λ_n is a path in Γ_n , so clearly $d_\Lambda(\sigma, \tau) \geq d(\sigma, \tau)$. We need to show that the distance between σ and τ in Γ_n can be realized by a path lying entirely within Λ_n .

We proceed by induction. For $n = 1$, we have of necessity $\sigma = e$, $\tau = e$, and $d_\Lambda(\sigma, \tau) = d(\sigma, \tau) = 0$. For $n > 1$, our aim will be to reduce to the following case: Suppose $\sigma(n) = n$ and $\tau(n) = n$ both. Then σ and τ can both be taken to lie within Γ_{n-1} and Λ_{n-1} , and by induction a shortest path within Γ_{n-1} , which is also a shortest path within Γ_n , can be taken to lie within $\Lambda_{n-1} \subset \Lambda_n$.

To reduce to this case, consider first the action of the duality map $(\sigma, \tau) \mapsto (\vec{\sigma}, \vec{\tau})$. This is an isometry of the Cayley graph Γ_n which exchanges the endpoints e and s of $\mathcal{I}_{e,s}$, and so it is also an isometry of Λ_n , so $d_\Lambda(\sigma, \tau) = d(\sigma, \tau)$ if and only if $d_\Lambda(\vec{\sigma}, \vec{\tau}) = d(\vec{\sigma}, \vec{\tau})$. The effect on M of taking simultaneous duals on σ and τ is to pull the first section of curve passing through $\bullet 1$ around to become $n\bullet$, relabeling $1\bullet$ as $\bullet 1$, $\bullet 2$ as $1\bullet$, and so forth.

Now, since the system of arcs in the lower half-plane is planar and noncrossing, there must be an innermost arc, an arc which encloses no other. By applying duality repeatedly, this innermost arc can be moved to the first position connecting $\bullet 1$ with $1\bullet$, and by applying duality twice more we may assume without loss of generality that $\tau(n) = n$. If $\sigma(n) = n$ also, then by induction we are done; otherwise we are halfway there.

Assume then that $\tau(n) = n$ but $\sigma(n) = i \neq n$. In M we thus have a curve which descends through point $\bullet n$, bends around immediately to pass through $n\bullet$, and then continues on to $\bullet i$. Let σ' be the product of permutations $(i\ n)\sigma$; then the system of meanders M' given by (σ', τ) has a circle through the points $\bullet n$ and

$n\bullet$, and the section of curve which used to pass down through $\bullet n$ instead doubles back up to eventually pass down through $\bullet i$. In particular,

- The permutation σ' corresponds to the noncrossing partition obtained by splitting out a singleton block $\{n\}$ from the noncrossing partition of σ , and is thus one step away from σ in Λ_n .
- The system of meanders M' has exactly one more component than M , so $\Pi_0(\sigma', \tau) = \Pi_0(\sigma, \tau) + 1$
- and thus $d(\sigma', \tau) = d(\sigma, \tau) - 1$.
- Since both $\tau(n) = n$ and $\sigma'(n) = n$, we can assume by induction that $d_\Lambda(\sigma', \tau) = d(\sigma', \tau)$.

It follows that a path realizing the length $d(\sigma, \tau)$ can be taken entirely within Λ_n , the first step of which is from σ to σ' . \square

Remark 3.4. David Savitt has proven independently [Sav06] that meanders correspond to diameters of the Hasse diagram of the lattice of noncrossing partitions.

3.3. Open questions.

- (1) This gives a definition of “meander” for any finite Coxeter group W . What can we say about these generalized meanders?
- (2) This suggests a generalization of the *meander determinant*. Do the generalized determinants also factor as nicely?
- (3) Does this give any hope for the meander enumeration problem?

REFERENCES

- [LZ93] S. K. Lando and A. K. Zvonkin. Plane and projective meanders. *Theoretical Computer Science*, 117(1–2):227–241, August 1993.
- [Sav06] David Savitt. Polynomials, meanders, and paths in the lattice of noncrossing partitions. 2006. <http://arxiv.org/abs/math.CO/0606169> (preprint).
- E-mail address:* `h.tracy@gmail.com`